



AN INFLUXDATA CASE STUDY

IBM Improves nmon Capabilities with InfluxDB and Grafana



SEPTEMBER 2022



Telegraf and client libraries facilitate data ingest that feeds visualizations

Company in brief

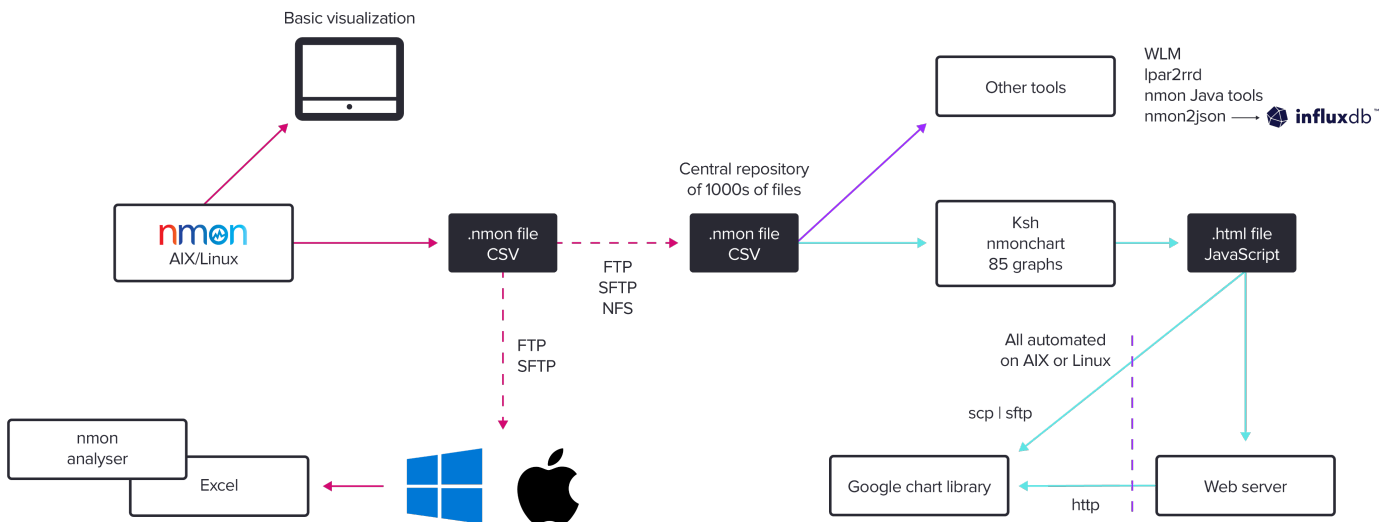
IBM is a multinational technology company headquartered in New York. IBM produces and sells a wide range of products and services, including computer hardware, middleware, and software, and provides hosting and consulting services in areas ranging from mainframe computers to nanotechnology.

Case overview

nmon is a tool for monitoring operating system performance. With over one million downloads, there is significant demand for the statistics the tool collects. However, Nigel Griffiths, the tool's creator, sought to update nmon to make it more useful for modern computing systems. He went about re-writing the tool to expand its capabilities and to automate processes, making it easier to use. He chose InfluxDB as the data storage system and Grafana for data visualization. These updates made the tool much more powerful and made it easier for users to visualize, analyze, and solve complex issues.

The business challenge

IBM engineer Nigel Griffiths built nmon in the 1990s to monitor operating system performance data for AIX. Griffiths later built an open source version for Linux, too. This monitoring tool can either display condensed data on screen or save it in a CSV file for later graphing and analysis. However, computer technology changed significantly since Griffiths first wrote nmon. In 2018, he sought to update the nmon tool to reflect those changes, make it more efficient, and add automation.



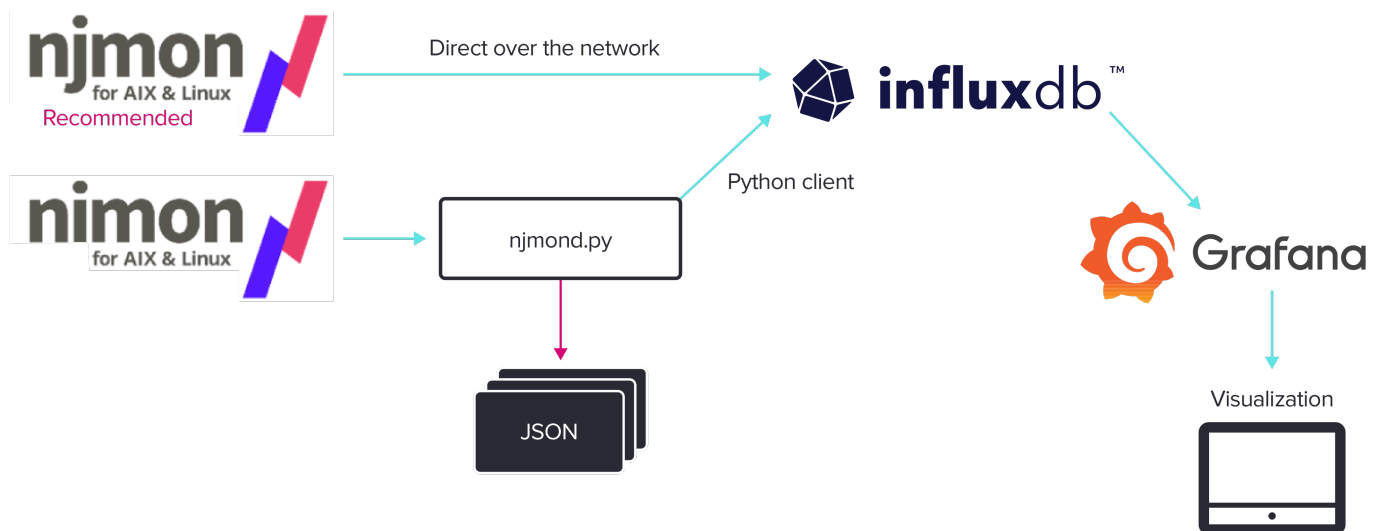
The technical challenge

Today's computers are orders of magnitude more powerful than machines in the 1990s, from processors, to RAM, to networking speed, to disk capacities, and more. The original nmon tool collected a limited number of system stats. On top of that, Griffiths describes the file format as quirky and non-standard. To modernize the tool for contemporary systems, Griffiths sought to increase the number of statistics collected, to standardize the file format, and to improve data storage and visualization.

The solution

Griffiths looked into the best solutions for his storage and graphing redesign. When he approached a trusted co-worker at one of IBM’s benchmark centers, the facilities that look after mainframe and power computers, they instantaneously responded with, “InfluxDB and Grafana. Don't even bother looking anywhere else.”

Ultimately, Griffiths followed this advice. He rewrote the entire tool to collect every possible operating system statistic. He updated the built-in file format for nmon and replaced his bespoke format with two standard formats, JSON and InfluxDB line protocol. This resulted in two different versions of the tool, njmon and nimon, respectively.



He also incorporated InfluxDB as the central database for storing all the statistics the tool processed. Using the JSON format allows nmon to push data to Elastic and Splunk. The JSON version also sends data to a central daemon written in Python, which uses the Python client library to push that data into InfluxDB. The line protocol version uses Telegraf to collect data and sends data to Prometheus.

Once data is in InfluxDB, nmon uses Grafana to visualize that data. The upgrade to Grafana allows users to visualize data in real time, instead of at the end of the hour or the end of the day, as was the case with the original version of nmon.

“

An nmon user told me, ‘This is the config file for Telegraf,’ and it was 10 lines long. I was shocked. First time I used Telegraf, I was just staggered about how clever that was in doing exactly what I wanted, and I thought it was going to take me months to try and work out how to solve that problem.

Nigel Griffiths, IBM

Results

These updates make nmon a much more robust and useful tool that’s on par with the needs of modern computer technology. nmon’s new capabilities enable users to explore and solve complex problems, like how to use InfluxDB and Grafana to spot trends in CPU usage in big production servers to alert on and avoid batch overruns.

What’s next

Griffiths continues to develop the nmon tool and is looking into ways to more easily document the names of the different statistics and metrics the tool collects. He’s also investigating the best ways to capture ad hoc statistics from big production servers.

About InfluxData

InfluxData is the creator of InfluxDB, the leading time series platform. We empower developers and organizations, such as Cisco, IBM, Lego, Siemens, and Tesla, to build transformative IoT, analytics and monitoring applications. Our technology is purpose-built to handle the massive volumes of time-stamped data produced by sensors, applications and computer infrastructure. Easy to start and scale, InfluxDB gives developers time to focus on the features and functionalities that give their apps a competitive edge. InfluxData is headquartered in San Francisco, with a workforce distributed throughout the U.S. and across Europe. For more information, visit influxdata.com and follow us [@InfluxDB](https://twitter.com/InfluxDB).



Try InfluxDB

Get InfluxDB

Contact us for a personalized demo influxdata.com/get-influxdb/